

# Testing Period for Migrating to the TOSS Operating System

UPDATE IN PROGRESS: The NAS Division is migrating its production environment from the SUSE Linux Enterprise Server (SLES) operating system to the Red Hat Enterprise Linux-based Tri-Lab Operating System Stack (TOSS), developed at the U.S. Department of Energy. The specific version used is TOSS 3.

This article provides some basic information to help you get started testing applications on the TOSS operating system. Updated information will be provided as testing and migration progresses.

## Front-End Systems

During the testing period, tfe1 will serve as the front-end system with the TOSS image. You can access tfe1 from a Pleiades or Lou front end (PFE or LFE). You can also add tfe1 to your `~/.ssh/config` file and SSH directly from your local workstation to tfe1.

Two PFEs, pfe26 and pfe27, have been migrated to TOSS. The remaining PFEs, pfe[20-25], are still running SLES and will be migrated to TOSS later this year.

As part of the transition, pfe26 and pfe27 have been removed from the load balancer. That is, connections to the PFEs established by running `ssh pfe` will only connect to pfe[20-25]. In order to connect with the TOSS PFEs, you must specify the number of the PFE. For example:

```
ssh pfe27
```

## Compute Nodes

Most of the compute node model types are available with the TOSS image:

### Model TOSS Available?

san	Yes
ivy	Yes
has	Yes
bro	Yes
sky_ele	Yes
cas_ait	Yes
rom_ait	Yes
san_gpu	No
sky_gpu	Yes
cas_gpu	Yes

You can verify that your node is running the TOSS image by running `cat /etc/os-release`. If the node is running TOSS, it will show **RHEL version 7.9**.

## Running PBS Jobs

### CPU Nodes Served by pbspl1

In the resource request of your PBS script or **qsub** command, add the attribute **:aoe=toss3**. For example:

A limited number of compute nodes of various models are already booted into the TOSS image and can be accessed via the **testing\_free** queue. Jobs submitted using the **testing\_free** queue will not be charged SBUs. Use the following **qstat** command on a PFE to find how many nodes are available before submitting to the queue.

```
pfe qstat au foo testing_free
pfe qsub q testing_free your_job_script
```

If you need more nodes with the TOSS image than are available in the **testing\_free** queue, you can use other queues (such as **devel**, **normal**, etc.), and PBS will provision nodes into the TOSS image. Be aware that PBS has to wait for nodes to become available and can only provision up to 50 nodes at a time, so it may take some time for your job to start.

## GPU Nodes Served by pbspl4

Two **sky\_gpu** and two **cas\_gpu** nodes have been booted to the TOSS image and can be accessed via the **testing\_free** queue on pbspl4.

Add the attribute **:aoe=toss3** in your resource request. For example, to request 1 **cas\_gpu** node with 4 GPU cards, do:

```
pbspl4 qsub I q testing_free lselectncpusngpusmemGBmodelcas_gpuaotoss3
```

To request a partial node with 12 CPU cores and 1 GPU card, do:

```
pbspl4 qsub I q testing_free lselectncpusngpusmemGBmodelcas_gpuaotoss3
```

Unlike nodes served by pbspl1, provisioning is not available for nodes served by pbspl4.

## Using Software Modules

On tfe1 and the compute nodes booted with a TOSS image, the default MODULEPATH includes the following:

```
usrsharemodulesmodulefiles
nasamodulefilestoss3
nasamodulefilesspackgcc
nasamodulefilespkgsrctoss3
```

Use the **module avail** command to find what software modules are added under these paths.

While some of these modules were built on a TOSS system, for example, **gcc/10.2**, **gcc/9.3**, and **pkgsrct/2020Q4**, most are SLES 12 versions and their modulefiles are simply copied from the **/nasa/modulefiles/sles12** module tree.

Please report any issues you encounter when using any of the modules, so that they can be removed from the default MODULEPATH and be rebuilt under the TOSS image.

There are many more modules under **/nasa/modulefiles/sles12** that have not been copied over. You can access these modulefiles by doing:

```
module use nasamodulefilessles12
```

```
module avail
module load name_of_a_sles12_module
```

Please report any success or failure in using the SLES 12 modules under the TOSS environment to help NAS staff determine what software modules need to be rebuilt under TOSS.

## Building Your Application

It is likely that your existing executables built under SLES 12 will also run under TOSS. If so, please check for both correctness and performance.

In the event that you need to build a fresh executable under TOSS, be aware that tfe1 is an 8-core Intel Xeon E5472 (Harpertown) system. It is best to compile on a compute node, so that Intel compiler flags such as `-fast` or `-xHost` will create an executable optimized for the intended architecture.

## Running X Applications on Compute Nodes

You can use the `-v` option of the `qsub` command from tfe1 to export the `DISPLAY` environment variable into the compute nodes. However, this is not recommended when you run `qsub` from a PFE, as other environment variables (such as `MODULEPATH`) from your SLES session will also be exported. You can set the `DISPLAY` variable manually after getting into the compute nodes.

```
pfe27
pfe27
pfe27 qsub I q testing_free lselectncpusmodelivyaotoss3
PBS r407i1n2 setenv DISPLAY pfe27
PBS r407i1n2 xclock
```

## Reporting Success or Issues

Please communicate with the NAS support team by sending email to [support@nas.nasa.gov](mailto:support@nas.nasa.gov).

---

Article ID: 668

Last updated: 13 Oct, 2021

Revision: 21

Filesystems & Software -> Software -> Testing Period for Migrating to the TOSS Operating System

<https://www.nas.nasa.gov/hecc/support/kb/entry/668/>